

What is claimed is:

1. A method for controlling scanning for computer viruses within a data processing network, comprising the steps of:

5        computing a set of hash values representing a set of resources;  
         in response to a requirement for a virus check, comparing the computed hash values to identify resources within said set of resources having matching hash values;  
         performing a virus scan for a first resource within said set of resources and, in response to the virus scan determining that the first resource is virus-free, recording a  
10        virus-free status for the first resource and the identified resources having hash values matching the hash value of the first resource.

2. The method of claim 1, wherein:

         the steps of comparing the hash values, performing a virus scan, and recording a  
15        virus-free status are performed at a first data processing system within the network for a set of resources distributed across a plurality of data processing systems within the network.

3. The method of claim 1, wherein:

20        the steps of computing hash values for a resource are performed at a data processing system within the network at which the resource is stored, the method further comprising sending the computed hash values to said first data processing system.

4. The method of claim 2, further comprising:

25        forwarding an indication of the virus-free status to a plurality of data processing systems at which a resource matching the first resource is stored.

5. The method of claim 2, further comprising:

         performing decontamination of the first resource at the first data processing  
30        system and forwarding a copy of the decontaminated first resource to data processing

systems, within the plurality of data processing systems, storing an identified resource having a hash values matching the hash value of the first resource.

6. A method for controlling scanning for computer viruses within a data processing network, comprising the steps of:

receiving a set of hash values derived by applying a secure hash function to each of a set of resources;

storing the set of hash values;

in response to a requirement for a virus check, comparing the computed hash values to identify resources within said set of resources having matching hash values;

performing a virus scan for a first resource within said set of resources and, in response to the virus scan determining that the first resource is virus-free, recording a virus-free status for the first resource and the identified resources having hash values matching the hash value of the first resource.

7. The method of claim 6, wherein storing the received hash values comprises updating a repository of hash values representing a set of resources stored at each of a plurality of data processing systems within a network.

8. The method of claim 7, wherein comparing the computed hash values comprises comparing hash values for resources stored at each of the plurality of data processing systems to identify resources replicated across a set of said data processing systems, performing the virus scan for one of the replicas of a resource, and forwarding the result of the virus scan to each of the set of data processing systems for recordal in relation to a respective replica resource.

9. The method of claim 8, further comprising:

using the identification of replicated resources to generate a report of the distribution of replicas of a resource.

10. A method for controlling performance of an operation within a data processing network, comprising the steps of:

computing a set of hash values representing a set of resources;

in response to a requirement for performance of the operation, comparing the

5 computed hash values to identify resources within said set of resources having matching hash values;

performing the operation in relation to a first resource within said set of resources and recording a result of the operation in association with the first resource and in association with identified resources having hash values matching the hash value of the  
10 first resource.

11. A method for controlling performance of an operation within a data processing network, comprising the steps of:

15 receiving a set of hash values derived by applying a secure hash function to each of a set of resources;

storing the set of hash values;

in response to a requirement for performance of the operation, comparing the computed hash values to identify resources within said set of resources having matching hash values;

20 performing the operation in relation to a first resource within said set of resources and recording a result of the operation in association with the first resource and identified resources having hash values matching the hash value of the first resource.

12. A data processing apparatus comprising:

25 a data processing unit;

a data storage unit;

a repository manager configured to store a set of hash values in at least one repository within the data storage unit, wherein the set of hash values are derived from a set of resources determined to be virus free; and

30 a virus scan coordinator for comparing the computed hash values to identify

resources having matching hash values, for controlling performance of a virus scan for a first resource, and for responding to said virus scan determining that the first resource is virus-free by controlling the repository manager to record a virus-free status in association with the first resource and resources having hash values matching the hash value of the first resource.

13. A data processing apparatus comprising:

a data processing unit;

a data storage unit;

a repository manager configured to store a set of hash values in at least one repository within the data storage unit, wherein the set of hash values are derived from a set of resources determined to be virus free; and

a coordinator for coordinating performance of an operation by comparing the computed hash values to identify resources having matching hash values, for controlling performance of the operation for a first resource, and for controlling the repository manager to record a result of the operation in association with the first resource and resources having hash values matching the hash value of the first resource.

14. The data processing apparatus of claim 13, further comprising:

a plurality of operator programs, each configured to respond to instructions from said coordinator to perform a respective operation in relation to the first resource.

15. The data processing apparatus of claim 14, wherein the plurality of operator programs comprises a plurality of virus scanning programs.

16. The data processing apparatus of claim 14, wherein the plurality of operator programs comprises a plurality of virus-decontaminator programs.

17. The data processing apparatus of claim 13, further comprising a report generator for generating a report of the distribution of resources having hash values matching the

hash value of the first resource.

18. A computer program product, comprising program code recorded on a recording medium, for controlling the performance of operations on a data processing system on which the program code executes, wherein the program code comprises:

a repository manager configured to store, in at least one repository, a set of hash values derived from a set of resources; and

a virus scan coordinator for comparing the computed hash values for the set of resources to identify resources having matching hash values, for controlling performance of a virus scan for a first resource, and for responding to a determination by said virus scan that the first resource is virus free by controlling the repository manager to record a virus-free status in respect of the first resource and resources having hash values matching the hash value of the first resource.

19. A method for controlling scanning for computer viruses within a data processing network, comprising the steps of:

receiving a set of hash values derived by applying a secure hash function to each of a set of resources;

storing the set of hash values;

in response to a requirement for a virus check, comparing the computed hash values to identify resources within said set of resources having matching hash values; performing a virus scan for a first resource within said set of resources and, in response to the virus scan determining that the first resource is virus-contaminated, recording a virus-contaminated status for the first resource and identified resources having hash values matching the hash value of the first resource.

20. The method of claim 19, wherein storing the set of hash values comprises updating a repository of hash values representing a set of resources stored at each of a plurality of data processing systems within a network.

21. The method of claim 20, wherein comparing the computed hash values comprises comparing hash values for resources stored at each of the plurality of data processing systems to identify resources replicated across a set of said data processing systems, performing the virus scan for one of the replicas of a resource, and forwarding the result of the virus scan to each of the set of data processing systems for recordal in relation to a respective replica resource.

22. A data processing apparatus comprising:

a data processing unit;

a data storage unit;

a repository manager configured to store a set of hash values in at least one repository within the data storage unit, wherein the set of hash values are derived from a set of resources determined to be virus free; and

a virus scan coordinator for comparing the computed hash values to identify resources having matching hash values, for controlling performance of a virus scan for a first resource, and for responding to said virus scan determining that the first resource is virus-contaminated by controlling the repository manager to record a virus-contaminated status in association with the first resource and resources having hash values matching the hash value of the first resource.

23. A computer program product, comprising program code recorded on a recording medium, for controlling the performance of operations on a data processing system on which the program code executes, wherein the program code comprises:

a repository manager configured to store, in at least one repository, a set of hash values derived from a set of resources; and

a virus scan coordinator for comparing the computed hash values for the set of resources to identify resources having matching hash values, for controlling performance of a virus scan for a first resource, and for responding to a determination by said virus scan that the first resource is virus-contaminated by controlling the repository manager to record a virus-contaminated status in respect of the first resource and resources having

hash values matching the hash value of the first resource.